

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (currently amended) A computer implemented method, comprising:
creating a ranking model that predicts a likelihood that a document will be selected by:
[[;]]
storing information associated with a plurality of prior searches;
determining a prior probability of selection based, at least in part, on the information associated with the prior searches, and generating the ranking model based, at least in part, on the prior probability of selection;
training the ranking model using a data set that includes approximately tens of millions of instances;
identifying documents relating to a search query;
scoring the documents based, at least in part, on the ranking model;
forming search results for the search query from the scored documents; and outputting the search results.
2. (canceled)
3. (previously presented) The method of claim 1, wherein the information associated with the prior searches includes, for each of a plurality of documents associated with

the prior searches, at least one of a position occupied by the document within prior search results, a score assigned to the document, or a number of documents listed above the document in the prior search results that were selected.

4. (original) The method of claim 1, wherein the creating a ranking model includes:
storing training data,
extracting features from the training data, and
generating conditions that include one or more of the extracted features.

5. (original) The method of claim 4, wherein the creating a ranking model further includes:
selecting one of the conditions as a candidate condition,
estimating a weight for the candidate condition,
forming a new rule from the candidate condition and the estimated weight,
comparing a likelihood of the training data between a current model with the new rule and the current model without the new rule, and
selectively adding the new rule to the current model based, at least in part, on a result of the comparison.

6. (previously presented) The method of claim 5, wherein the selecting one of the conditions as a candidate condition includes at least one of:
creating the candidate condition from combinations of features or complements of

features in the training data,

randomly selecting one of the conditions as the candidate condition,

selecting one of the conditions that includes a single one of the features as the candidate condition, or

augmenting one of the conditions by adding one or more features to the one condition to form the candidate condition.

7. (original) The method of claim 5, wherein the estimating a weight includes determining a weight that maximizes a likelihood of the training data given the model.

8. (previously presented) The method of claim 5, wherein the selectively adding the new rule to the current model includes adding the new rule to the current model when a likelihood of the training data occurring when the current model includes the new rule is greater than when the current model does not include the new rule.

9. (original) The method of claim 5, wherein the selectively adding the new rule to the current model further includes:

associating a cost with each of the conditions, and

determining whether to add the new rule to the current model based, at least in part, on the cost associated with the candidate condition.

10. (original) The method of claim 5, further comprising:

performing a number of iterations including estimating the weight, forming the new rule, and comparing the likelihood of the training data.

11. (original) The method of claim 1, wherein the data set also includes approximately millions of features.

12. (original) The method of claim 1, wherein the scoring the documents includes: forming an instance that corresponds to the search query and one of the documents, extracting features associated with the instance, identifying rules in the ranking model that apply based, at least in part, on the extracted features, each of the identified rules including a weight, and combining the weights of the identified rules with a prior probability of selection corresponding to the instance to generate a score for the one document.

13. (original) The method of claim 12, wherein the instance includes user information corresponding to a user who provided the search query, query data corresponding to the search query, and document information corresponding to the one document.

14. (original) The method of claim 1, wherein the scoring the documents includes: determining a prior probability of selection corresponding to the search query and one of the documents, and generating a score for the one document based, at least in part, on the determined prior

probability of selection.

15. (original) The method of claim 14, wherein the generating a score for the one document includes using the determined prior probability of selection as one of a plurality of factors in determining the score for the one document.

16. (previously presented) A system implemented within one or more computer devices, comprising:

means for receiving a search query;

means for identifying documents relating to the search query;

means for ranking the documents based, at least in part, on a ranking model trained on a large data set that includes approximately millions of features, the means for ranking includes:

means for determining a prior probability of selection corresponding to the search query and one of the identified documents, and

means for determining a rank for the one document based, at least in part, on the determined prior probability of selection;

means for forming search results for the search query from the ranked documents; and
means for outputting the search results.

17. (currently amended) A system, comprising:

a repository configured to store information corresponding to a plurality of prior searches; and

a server configured to:

receive a search query from a user,

identify documents corresponding to the search query,

rank the identified documents based, at least in part, on a ranking model that includes rules that maximize a likelihood of the repository, when ranking the identified documents, the server is configured to: [[;]]

determine a prior probability of selection corresponding to the search query and one of the identified documents, and

determine a rank for the one document based, at least in part, on the determined prior probability of selection, and output the ranked documents.

18. (original) The system of claim 17, wherein the information in the repository includes a plurality of instances that include user information, query data, and document information corresponding to the plurality of prior searches.

19. (original) The system of claim 18, wherein the repository is further configured to store a plurality of features associated with the instances.

20. (original) The system of claim 19, wherein when ranking the documents, the server is configured to:

identify one of the instances that corresponds to the search query and one of the identified

documents,

determine features associated with the identified instance,
identify rules in the ranking model that apply based, at least in part, on the determined
features, each of the identified rules including a weight, and
combine the weights of the identified rules with a prior probability of selection
corresponding to the identified instance to determine a rank for the one document.

21. (original) The system of claim 20, wherein the user information of the identified
instance includes information corresponding to the user who provided the search query, the
query data of the identified instance includes information corresponding to the search query, and
the document information of the identified instance includes information corresponding to the
one document.

22. (canceled)

23. (previously presented) The system of claim 17, wherein when determining a rank
for the one document, the server is configured to use the determined prior probability of selection
as one of a plurality of factors in determining the rank for the one document.

24. (original) The system of claim 17, wherein the repository stores approximately
tens of millions of instances and approximately millions of features associated with the plurality
of prior searches.

25-39. (canceled)

40. (previously presented) A method, comprising:
receiving a search query;
identifying documents relating to the search query;
determining prior probabilities of selecting each of the documents, where the prior probability of selecting one of the documents is determined based, at least in part, on data regarding at least one of a position of the document within search results, a prior score assigned to the document, or a number of documents above the document in the search results that were selected;
determining a score for each of the documents based, at least in part, on the prior probability of selecting the document;
generating search results for the search query from the scored documents, and outputting the search results.

41-46. (canceled)

47. (previously presented) A method, comprising:
creating a ranking model that predicts a likelihood that a document will be selected by:
storing information associated with a plurality of prior searches,
determining a prior probability of selection based, at least in part, on the

information associated with the prior searches, and
generating the ranking model based, at least in part, on the prior probability of
selection;
identifying documents relating to a search query;
scoring the documents based, at least in part, on the ranking model;
forming search results for the search query from the scored documents; and
outputting the search results.

48. (previously presented) A method, comprising:
receiving a search query;
identifying documents relating to the search query;
determining a prior probability of selecting one of the documents, the prior probability of
selecting the one document is determined based, at least in part, on data regarding at least one of
a position of the one document within search results, a prior score assigned to the one document,
or a number of documents above the one document in the search results that were selected;
determining a score for the one document based, at least in part, on the prior probability
of selecting the one document;
generating a list of search results that includes the one document based on the determined
score; and
outputting the list of search results.